



## Leap to Petascale

Developing Developer Tools  
TotalView and Blue Gene/Q

May 23 2012

Ed Hinkel,  
Sr Sales Engineer  
Rogue Wave Software

# Agenda

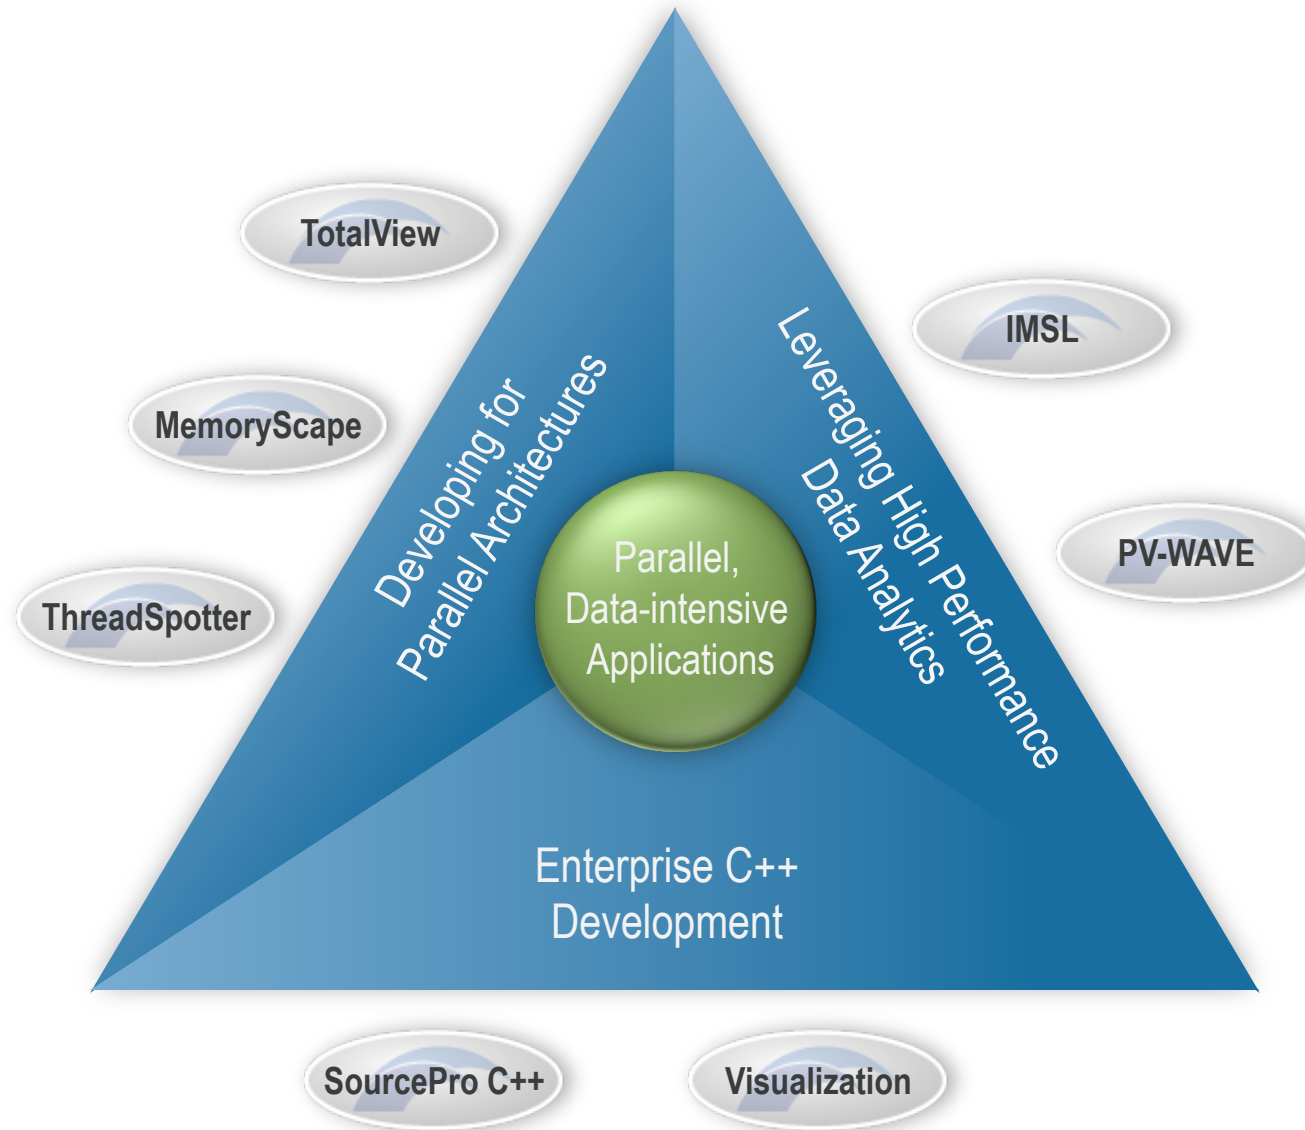
---

- **Who is Rogue Wave?**
- **Early Blue Gene Days with TotalView**
- **Blue Gene/Q Advancements**
- **Techniques for Debugging Challenges**
- **What's New with TotalView**

# Who is Rogue Wave Software?

## Solution Portfolio

---



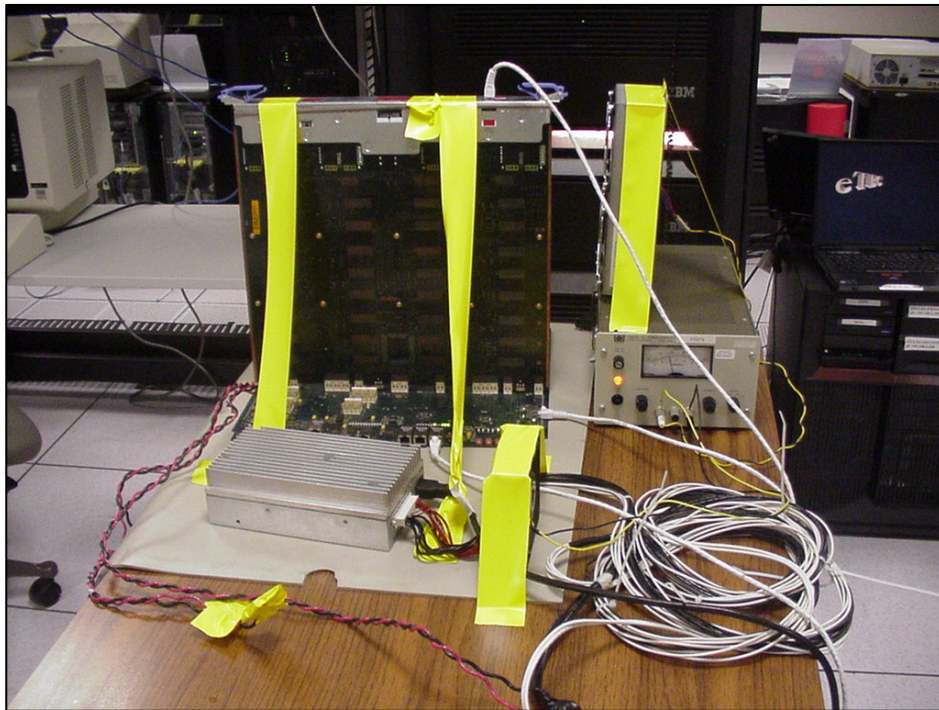
---

# Early Blue Gene Days with TotalView

# TotalView Blue Gene Support

---

- **TotalView involvement started in 2003 on BG/L**



IBM/TV BG/L development system

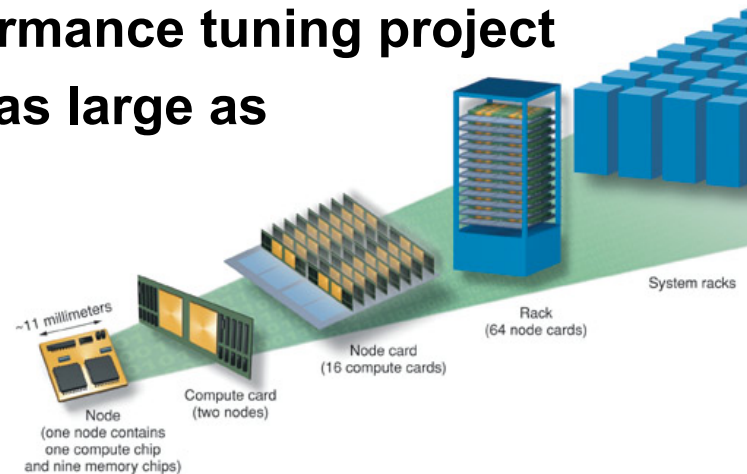
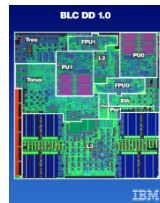


**Gotta love that yellow duct tape!**

# TotalView Blue Gene/L Support

---

- **Support for Blue Gene/L since 2005**
- **Debugging interfaces developed via close collaboration with IBM (CIOD)**
- **Used on DOE/NNSA/LLNL's Blue Gene/L system containing 212 K cores**
  - **Heap memory debugging support added**
  - **Blue Gene/L scaling and performance tuning project**
  - **TotalView has debugged jobs as large as 32,768 processes**



## Blue Gene/L work facilitated Blue Gene/P support

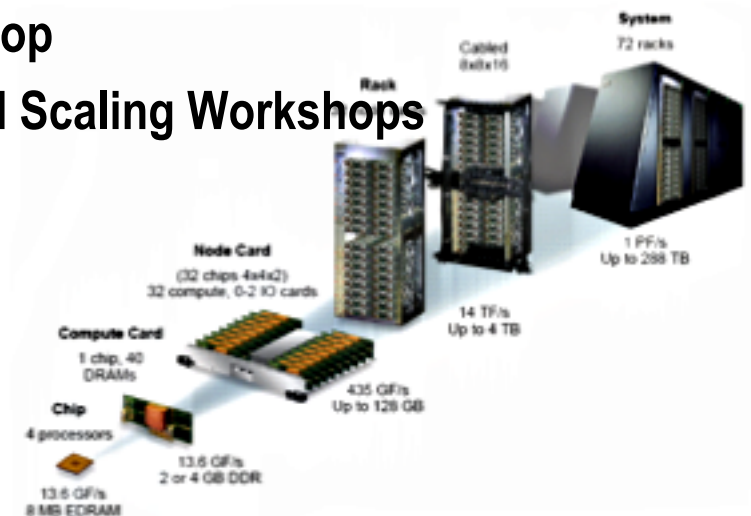




# TotalView Blue Gene/P Support

---

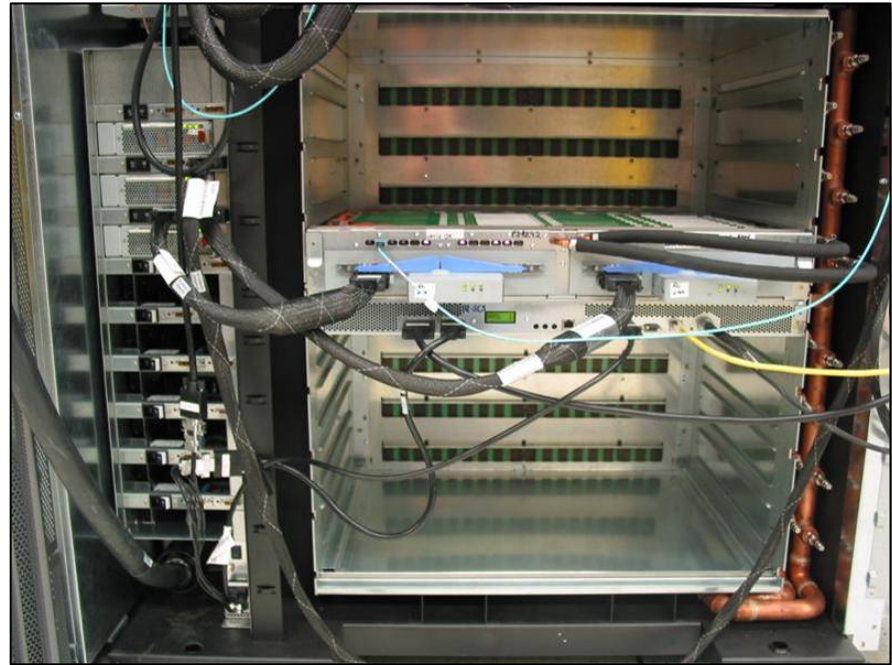
- Continued close collaboration with IBM
- Currently running on several BG/P installations in Germany, France, the UK, and the US.
- Support for shared libraries, threads, and OpenMP
- TotalView has debugged jobs as large as 32,768
- Active workshop participation through the development
  - ANL's ALCF INCITE Performance Workshop
  - Jülich's Blue Gene/P Porting, Tuning, and Scaling Workshops



# TotalView Blue Gene/Q Support

---

- **Porting TotalView began in June 2011**
- **Access to Q32 at IBM began in August**
- **Basic debugging operations in October**
- **Used in Synthetic Workload Testing in December**
- **Fully functional in March 2012**

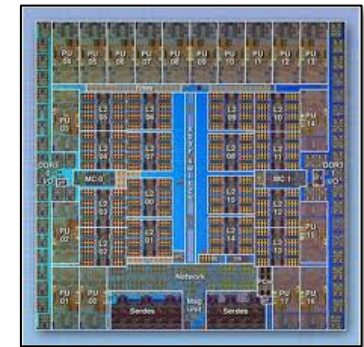


**IBM's Q32**

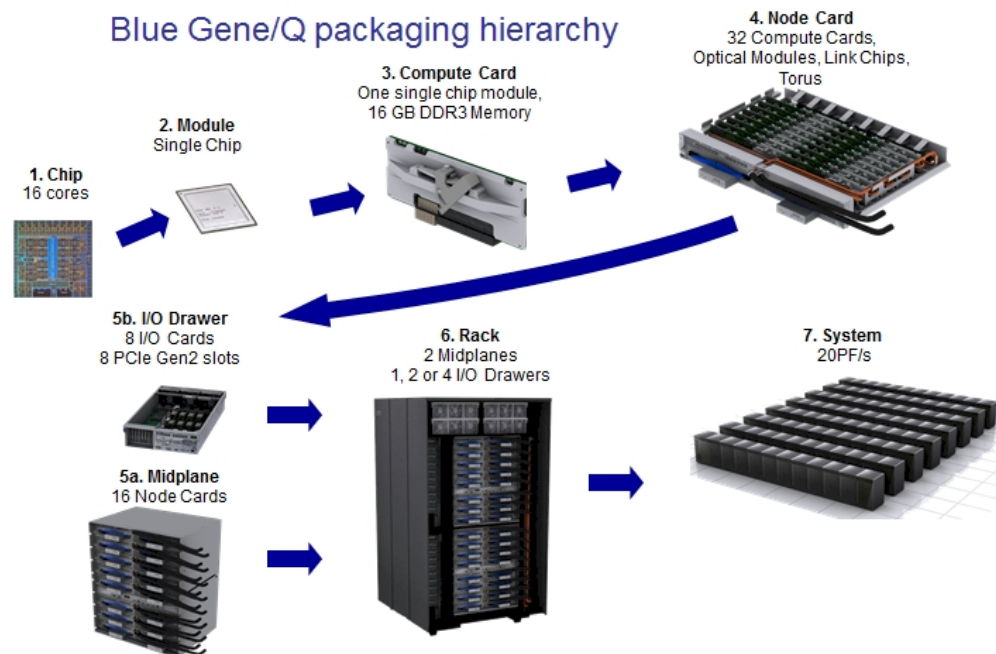


# TotalView Blue Gene/Q Support (cont)

- Thanks to the ongoing collaboration with IBM and the BG Kernel Team, early access versions of TotalView for BG/Q is available
- Argonne National Laboratory
- Lawrence Livermore National Laboratory



Blue Gene/Q packaging hierarchy



---

# **Blue Gene/Q Advancements with TotalView**

# TotalView on BG/Q Support

---

- **BG/Q TotalView is as functional as BG/P TotalView**
  - MPI, OpenMP, pthreads, hybrid MPI+threads
  - C, C++, Fortran, assembler; IBM and GNU compilers
  - Basics: source code, variables, breakpoints, watchpoints, stacks, single stepping, read/write memory/registers, conditional breakpoints, etc.
  - Memory debugging, message queues, binary core files, etc.
- **PLUS, features unique to BG/Q TotalView**
  - QPX (floating point) instruction set and register model
  - Fast compiled conditional breakpoints and watchpoints
  - Asynchronous thread control
- **Working with IBM on debugging interfaces for TM/SE regions**
  - TM == transactional memory; SE == speculative execution

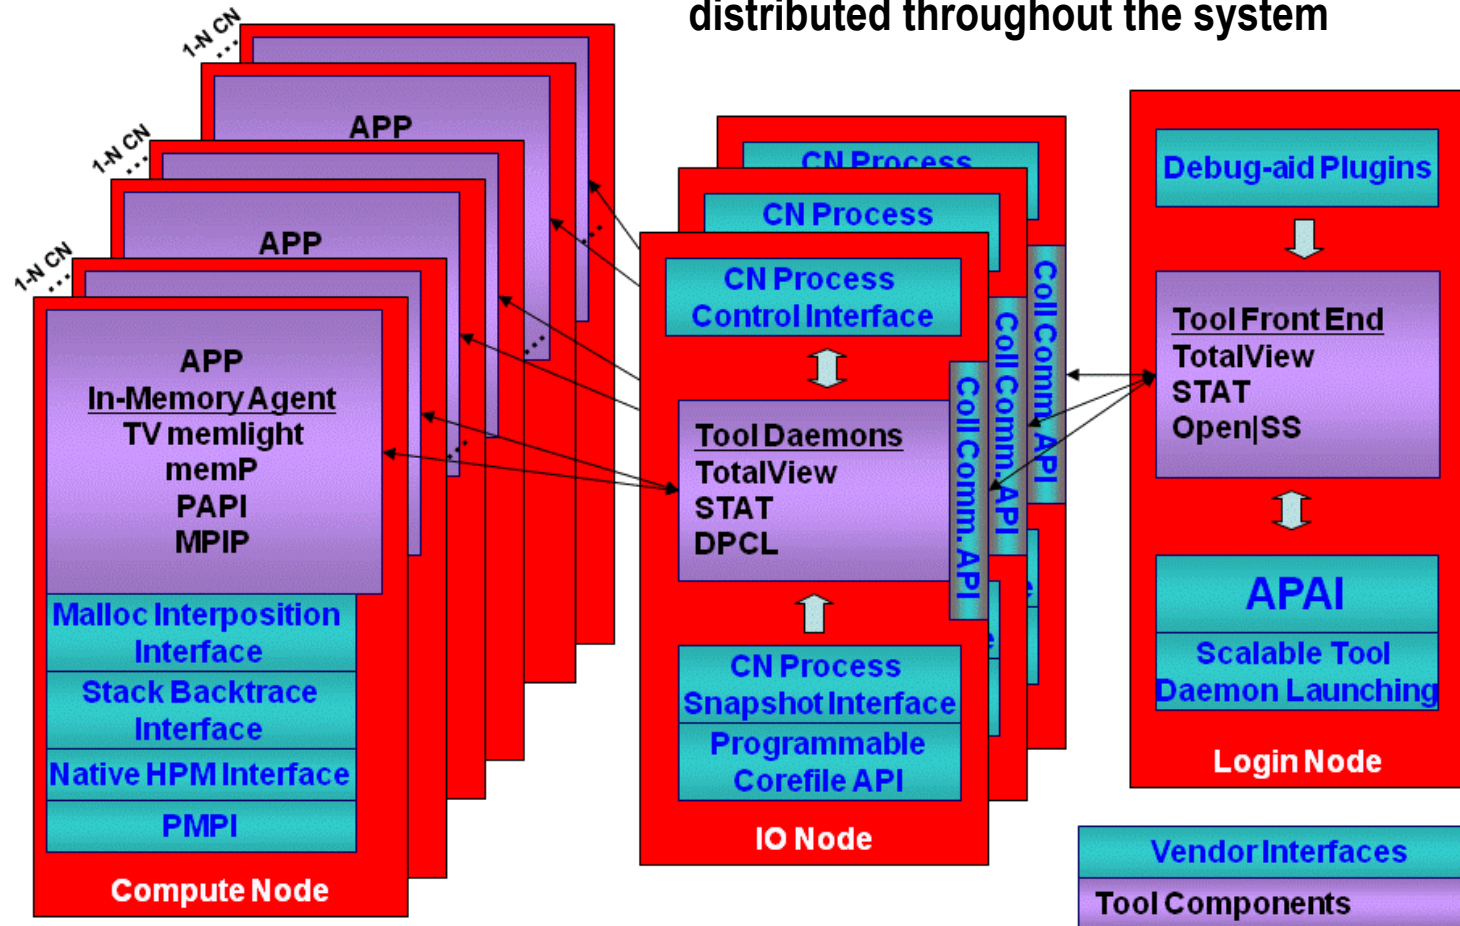
# Advanced TotalView Features on BG/Q

---

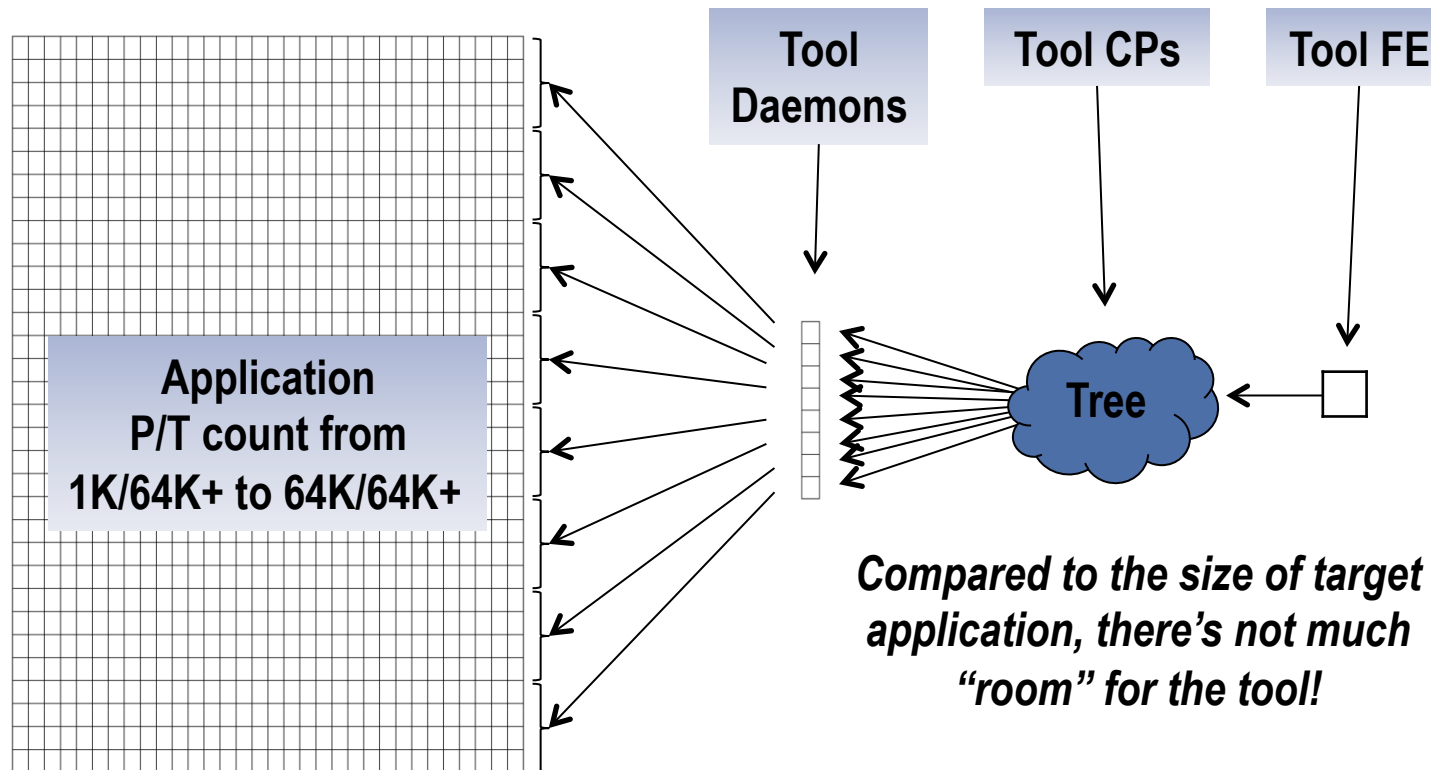
- **Asynchronous thread control**
  - A feature on Linux, and other TotalView platforms, ported to BG/Q
  - Allows you to individually control the execution of threads
  - Run and halt individual threads
  - Single-step a group of threads in lockstep
  - Hold and release the execution of individual threads
  - Create stop-thread and thread barrier breakpoints
- **Fast compiled conditional breakpoints and watchpoints**
  - A feature on AIX and other TotalView platforms, ported to BG/Q
  - Conditional breakpoints and watchpoints execute in as little as 7  $\mu$ secs
  - Conditional expressions are compiled and dynamically patched into the process
  - Evaluation is performed by the triggering thread, in parallel

# Blue Gene Code Development Tools Interface (CDTI)

Hierarchical infrastructure components are distributed throughout the system



# Tool Challenges



**One rack of BG/Q:  
1K CNs, 16K Cores,  
64K HW Threads**

**A "generous"  
128:1 CN:ION  
ratio: 8 IONs**

**A "beefy" FEN  
P7, 3 GHz+, 32 GB+**



# Overcoming High CN:ION Ratios

---

- On BG/Q, at a given ratio, on *each* IO node, tool daemons may be responsible for up to

CN:ION	Processes	Threads
64:1	4,096	20,480
128:1	8,192	40,960
256:1	16,384	81,920
512:1	32,768	163,840

- But each IO node has
  - 1.6 GHz A2 17 core processor (not too swift)
  - 16 GB (limited physical memory)

# Where to put the “weight” of the debugger?

---

- **Most of the “weight” of the debugger is in the symbol table**
- **Real-world applications are huge and complex**
- **A recently analyzed mission critical application revealed**
  - 1.5 million function definitions
  - 16 million line number definitions
  - DWARF symbol information in excess of 2 GB
  - 100s or 1000 of shared libraries
- **You don’t want to be big in the back end!**
- **And nothing too compute intensive either**

# TotalView's Architecture

---

- **Extremely lightweight back-end daemon processes**
  - Small footprint plus a few hundred bytes per CN process or thread
  - Each daemon can handle thousands of processes and threads
  - The daemons do not store the symbol table!
- **The “weight” of the debugger is on the front-end node**
- **Symbol tables are indexed and stored on the FEN**
  - Debugger has exactly *one* copy of the symbol table for each image file
  - Symbol tables are *shared* across all processes and thread
  - Aggregate memory consumption is minimal
- **Well suited to Blue Gene!**

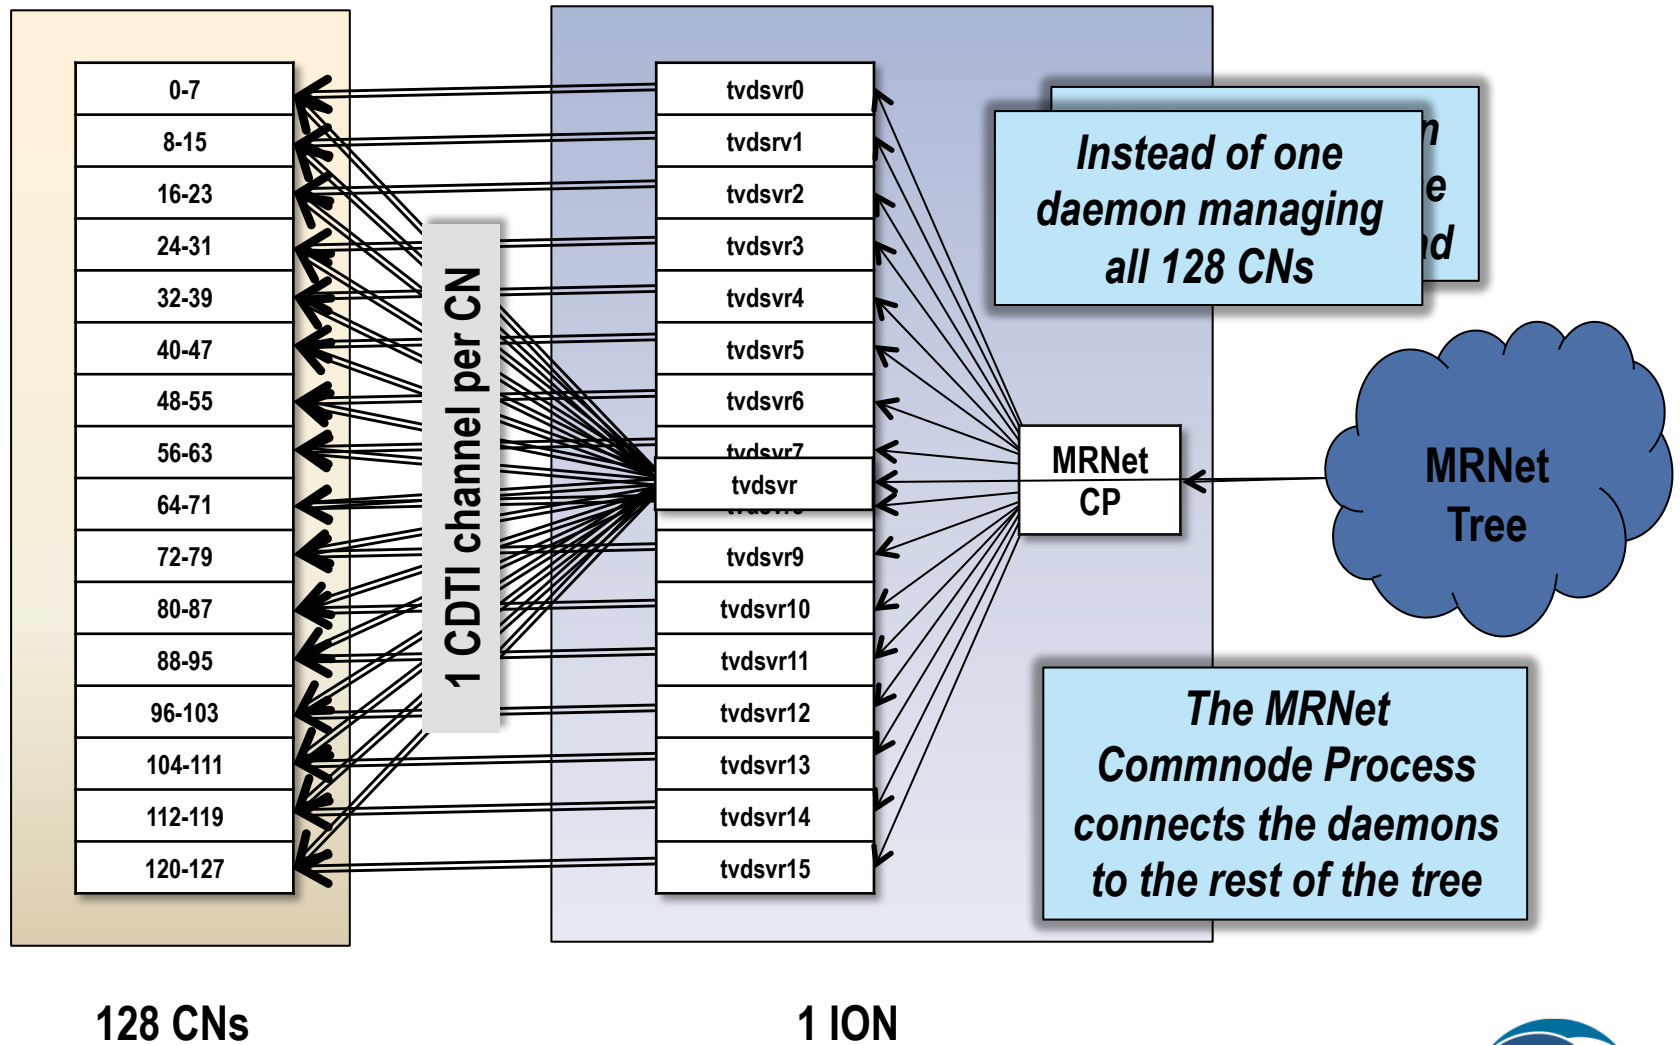
# There's still the high P/T count per IO node problem

---

- **Process and threads counts per IO node are still high!**
- **What to do about that?**
- **“Divide and conquer”**
  - Place a small number of daemons on the ION
  - We do have 17 cores we can use
- **Unlike CIOD on BG/L&P, CDTI on BG/Q can operate in parallel**
  - There's one CDTI debug channel per compute node



# Solution: TotalView/MRNet Trees on the IO Nodes



# TotalView on BG/Q Support

---

- **BG/Q TotalView is as functional as BG/P TotalView**
  - MPI, OpenMP, pthreads, hybrid MPI+threads
  - C, C++, Fortran, assembler; IBM and GNU compilers
  - Basics: source code, variables, breakpoints, watchpoints, stacks, single stepping, read/write memory/registers, conditional breakpoints, static/dynamic executables, etc.
  - Memory debugging, message queues, binary core files, etc.
- **PLUS, advanced BG/Q TotalView features**
  - QPX (floating point) instruction set and register model
  - Fast compiled conditional breakpoints and watchpoints
  - Asynchronous thread control
- **Working with IBM on debug interfaces for TM/SE regions**
  - TM == transactional memory; SE == speculative execution



# Advanced BG/Q TotalView Features

---

- **Asynchronous thread control**
  - A TotalView feature on Linux and other platforms, ported to BG/Q
  - Allows you to individually control the execution of threads
  - Run and halt individual threads
  - Single-step a group of threads in lockstep
  - Hold and release the execution of individual threads
  - Create stop-thread and thread barrier breakpoints
- **Fast compiled conditional breakpoints and watchpoints**
  - A TotalView feature on AIX and other platforms, ported to BG/Q
  - Conditional breakpoints and watchpoints execute in as little as 7  $\mu$ secs
  - Conditional expressions are compiled and dynamically patched into the process
  - Evaluation is performed in parallel by the triggering thread

# TotalView Availability

---

- **TotalView on Blue Gene/Q Today**
  - LLNL has it up and running on rzuseq, and is using it to debug applications.
  - IBM is using it internally for debugging and testing.
  - It's installed on IBM's Blue Gene On Demand Center Q32 (if anyone has access to that system).
- **TotalView On Blue Gene at Argonne**
  - 1024 Tokens (BG/P)
  - Research license is available with 65,536 tokens

# TotalView on VEAS!

The screenshot displays a VNC session with the following components:

- Terminal Window (Top Left):** Shows the output of a benchmark script. It includes network statistics (collisions, tx/rx bytes) and a list of files/directories used in the benchmark.
- File Manager (Bottom Left):** Displays the contents of the directory `/home/jdelsign/tw249_tests/fcbwbgq/benchmark`. It shows a list of files and directories with their permissions, owners, and timestamps.
- TotalView GUI (Top Right):** Shows the TotalView 8x.9.9-10 interface. The main window displays a table of processes, including `runjob<evaltest.stexe>.0`. The status of the process is `Stopped`.
- Code Editor (Bottom Right):** Displays the source code of the `main` function in `evaltest.cxx`. The code includes MPI initialization and a loop for processing arguments.

The code in the editor is as follows:

```

376 extern "C" void init_l2_atomic();
377 extern "C" int bgpm_start();
378 extern "C" void bgpm_finish(int);
379
380 int main(int argc, char** argv)
381 {
382     #ifdef USE_MPI
383         MPI_Init(&argc, &argv);
384
385         int myrank, numprocs;
386         MPI_Comm_rank(MPI_COMM_WORLD, &myrank);
387         MPI_Comm_size(MPI_COMM_WORLD, &numprocs);
388         MPI_Barrier(MPI_COMM_WORLD);
389     #endif
390
391     /* Process argv */
392
393     /* Test signal stack frame unw
394     establish sigquit_siginfo_handler (sigquit_siginfo_handler);
395     kill (getpid(), SIGQUIT);
396     establish sigquit_signal_handler (sigquit_signal_handler);
397     kill (getpid(), SIGQUIT);
398     signal (SIGQUIT, sigquit_signal_handler);
  
```

# TotalView on VEAS!

The screenshot displays the TotalView 8X.9.9-9 interface, which is used for debugging and monitoring applications. The interface is divided into several panels:

- Top-Left Panel:** A file browser showing a directory structure with files like `mispatch.c`, `mispatch.stex`, `omp_directive.c`, `omp_directive_repro.sh`, `omp_directive_repro.tvd`, `patchspace.lds`, `patchspace.s`, `qdemo.c`, `qdemo.dyx`, `qdemo.stex`, `run_snapbug.sh`, `sigtraptest.c`, `sigtraptest.cxx`, `totalview.sh`, and `totalview.sh`.
- Top-Right Panel:** A status window showing a table of threads. The table has columns for ID, Rank, Host, Status, and Description. The status is set to "Running (1 active threads)".
- Bottom-Left Panel:** A terminal window showing the output of a command: `runjob<evaltest.stex>.0 - 2.1`. The output shows the execution of a script that sets up a TotalView environment and runs a benchmark.
- Bottom-Right Panel:** A stack trace and a list of action points. The stack trace shows the function `main` in `evaltest.cxx`. The action points table shows the execution of various functions, including `main`, `generic_start_main`, `_libc_start_main`, `process_argv`, `kill`, `signal`, `init_l2_atomic`, and `USE_BQPM`.

---

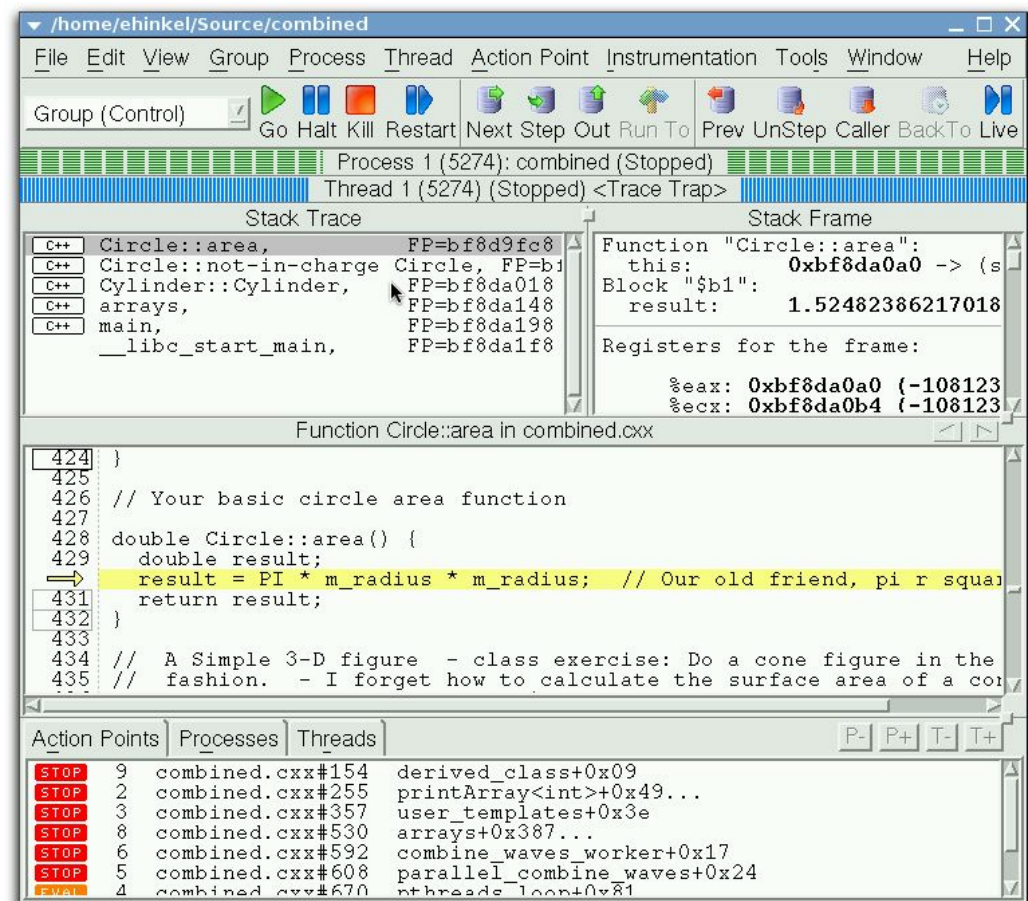
# Techniques for Debugging Challenges



# What is TotalView?

## A comprehensive debugging solution for demanding parallel and multi-core applications

- **Wide compiler & platform support**
  - C, C++, Fortran 77 & 90, UPC
  - Unix, Linux, OS X
- **Handles Concurrency**
  - Multi-threaded Debugging
  - Multi-process Debugging
- **Integrated Memory Debugging**
- **Reverse Debugging available**
- **Supports Multiple Usage Models**
  - Powerful and Easy GUI – Highly Graphical
  - CLI for Scripting
  - Long Distance Remote Debugging
  - Unattended Batch Debugging





# Debugging Complex Codes

---

- **Mechanize**
- **Minimize**
- **Visualize**
- **... and Don't Forget the Memory**

# Mechanize

---

## Extended Automation Capabilities



# Automated Debugging

---

## TVscript

- **Non-Interactive Batch Debugging –**
  - Work in the “main” batch queue
  - Don’t have to baby-sit job waiting on it to run
  - Use scripting to perform checks that would be tedious to do by hand
  - Verification through automated processes (nightly build and test)

## TTF and C++View

- **Automatic Transformation of Data –**
  - Simplify interactive (and scripted) debugging
  - Perform validation/sanity checking of large datasets
  - Comparative debugging
  - Allows you to focus on troubleshooting your program

# Non-Interactive Batch Debugging with TVscript

---

- Run multiple debugging sessions without the need for recompiling, unlike with printf
- TVscript syntax:  
`tvscript [ options ] [ filename ] [ -a program_args ]`
- More complex actions-to-events are possible, utilizing TCL within a CLI file
- TVscript lets you define what events to act on, and what actions to take

TVscript uses a simple, Event/Action interface	
<b>Typical Events</b> <ul style="list-style-type: none"><li>• Action_point</li><li>• Any_memory_event</li><li>• Guard_corruption error</li></ul>	<b>Typical Actions</b> <ul style="list-style-type: none"><li>• Display_backtrace [-level <i>level-num</i>]</li><li>• List_leaks</li><li>• Save_memory</li><li>• Print [-slice {<i>slice_exp</i>} {<i>variable</i>   <i>exp</i>}</li></ul>

# Unattended Debugging with Tvscrip

---

## Example

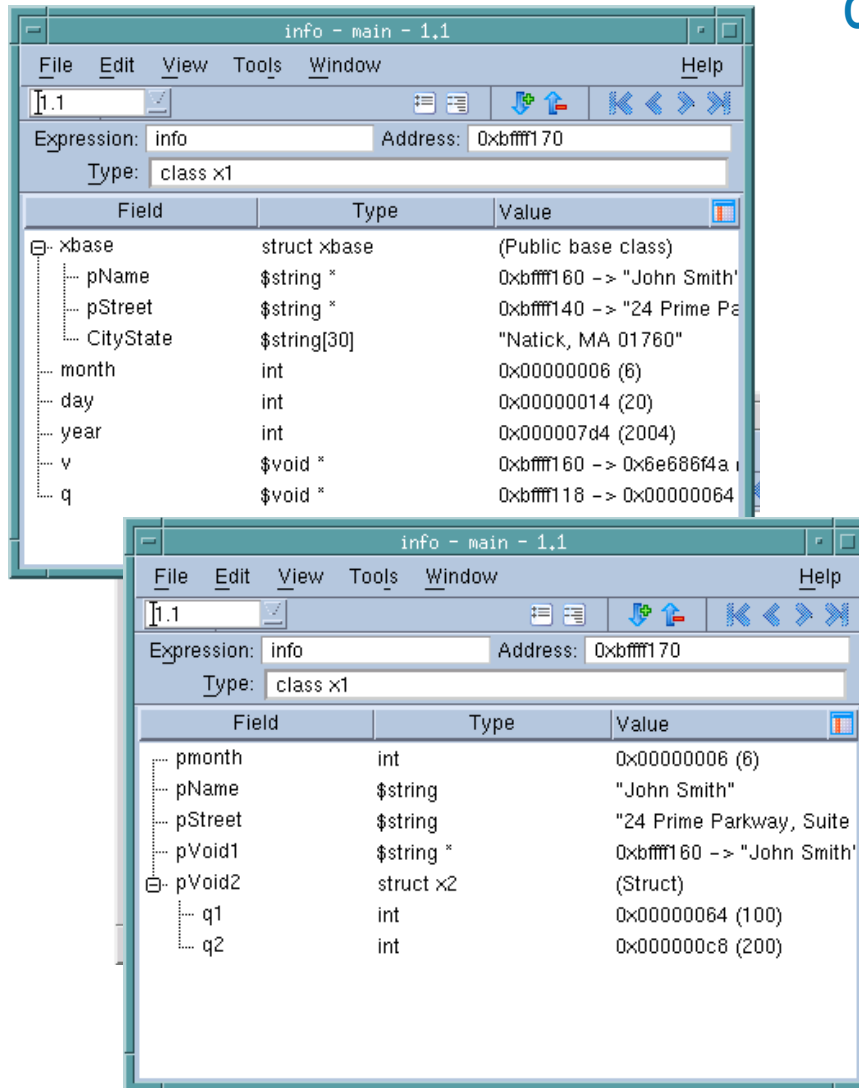
The following tells tvscript to report the contents of the *foreign\_addr* structure each time the program gets to line 85

**-create\_actionpoint "#85=>print foreign\_addr"**

Typical output sample with tvscript:

```
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
! Print
!
! Process:
!   ./server (Debugger Process ID:  1, System ID:  12110)
! Thread:
!   Debugger ID:  1.1, System ID:  3083946656
! Time Stamp:
!   06-26-2008 14:04:09
! Triggered from event:
!   actionpoint
! Results:
!   foreign_addr = {
!       sin_family = 0x0002 (2)
!       sin_port = 0x1fb6 (8118)
!       sin_addr = {
!           s_addr = 0x6658a8c0 (1717086400)
!       }
!       sin_zero = ""
!   }
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
```

# Creating Type Transformations



## Customize your own Transformations

In \$HOME/.tvdrc:

```
::TV::TTF::RTF::build_struct_transform {
    name {^class x1$}
    members {
        { pmonth { month }}
        { pName { xbase upcast { * pName }}}
        { pStreet { xbase upcast { * pStreet }}}
        { pVoid1 { "$string *" cast v }}
        { pVoid2 { * { "class x2 *" cast q }}}
    }
}
```

Meta Language:

```
{member}
{* expr}
{expr . Expr}
{expr -> expr}
{datatype case expr}
{baseclass upcast expr}
```



# C++View

- C++View is an easy way to customize TotalView's display of object data.
- How does it work?
  - User writes short display functions within their program
  - TotalView uses these functions to simplify the display of data when the user explores their data within that program
  - C++View transforms are easy to define
  - Great for collaborative codes (transforms can be distributed with the program)
- Benefit: Easier for scientists and developers to work with complex applications

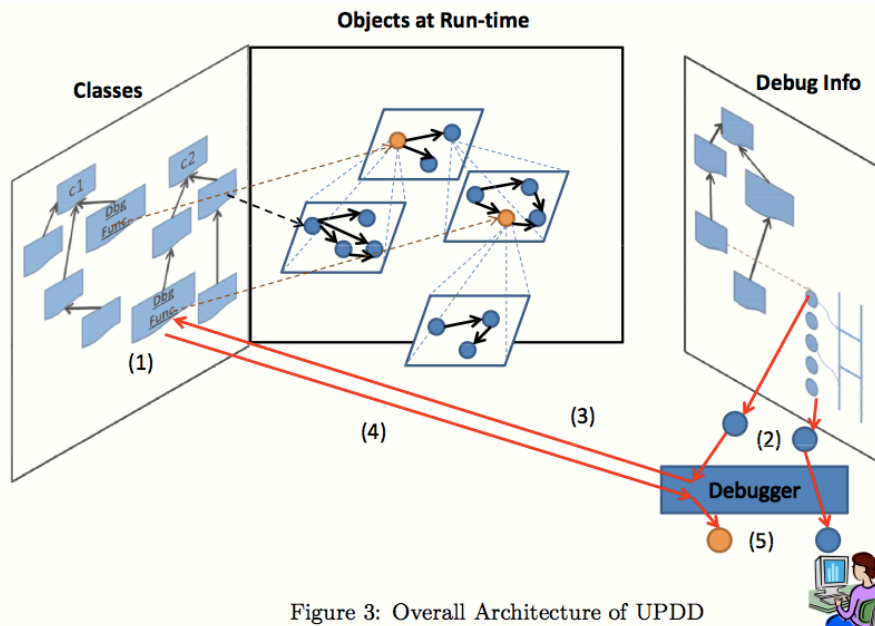


Figure 3: Overall Architecture of UPDD

Developers can now write display and analysis functions for their C++ classes that are invoked whenever an object is inspected interactively in the debugger.

- 
- The screenshot displays two windows from the Visual Studio IDE. The left window, titled '.\milestone\_example', shows the source code of the 'main' function in 'milestone\_example.cxx'. A breakpoint is set at line 45, which is highlighted. The code involves a loop over a vector of vectors of doubles, printing information about each element. The right window, titled 'triangle - main - 1.1', shows the evaluation of the 'triangle' expression. The expression is 'triangle', located at address '0xbfb05754'. Its type is 'struct std::vector<std::vector<double, std::allocator<double> >, std::allocator<std::vector<double, std::allocator<double> > > >'. The window displays a table of the vector's contents:
- | Field         | Type      | Value    |
|---------------|-----------|----------|
| at(0).front() | \$string  | "empty!" |
| at(1).front() | double[1] | (Array)  |
| [0]           | double    | 0        |
| at(2).front() | double[2] | (Array)  |
| [0]           | double    | 0        |
| [1]           | double    | 2        |
| at(3).front() | \$string  | "empty!" |
| at(4).front() | \$string  | "empty!" |
| at(5).front() | \$string  | "empty!" |

# Minimize

---

Reduce the Scope of Effort



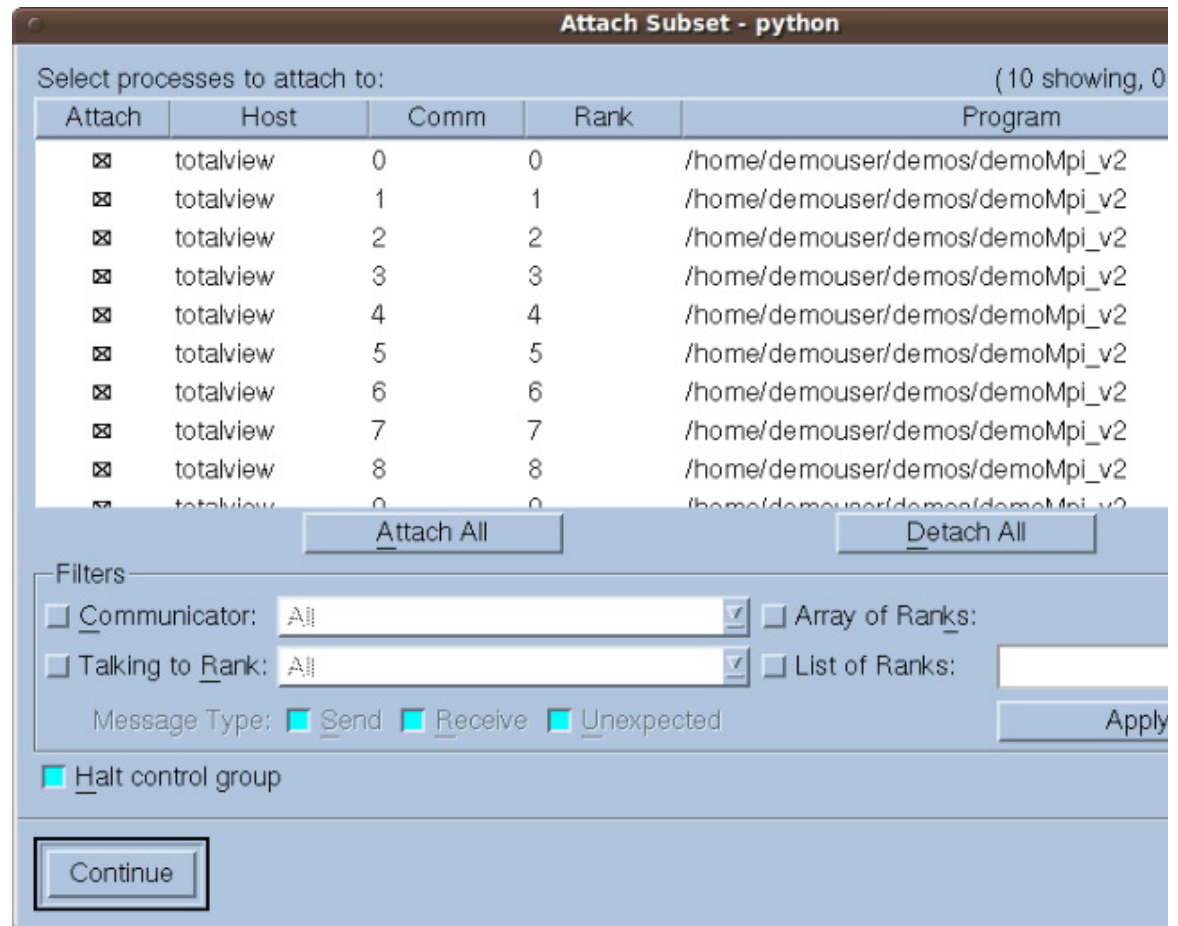
---

# Subset Debugging With TotalView

# Subset Attach

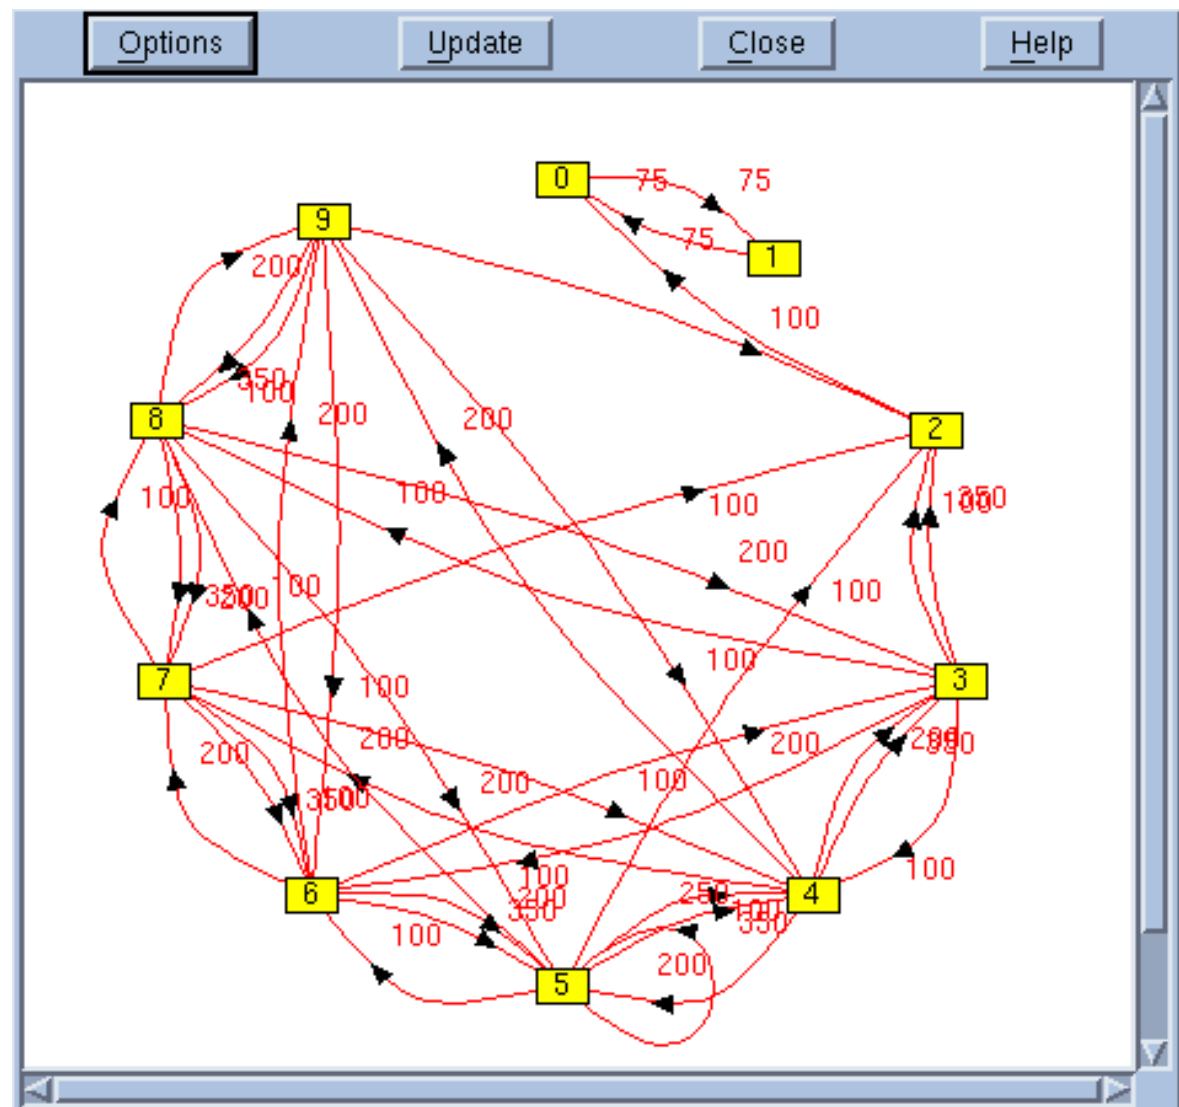
You need not be attached to the entire job

- You can be attached to different subsets at different times through the run
- You can attach to a subset, run till you see trouble and then 'fan out' to look at more processes if necessary.
- This greatly reduces overhead
- It also reduces license size requirements



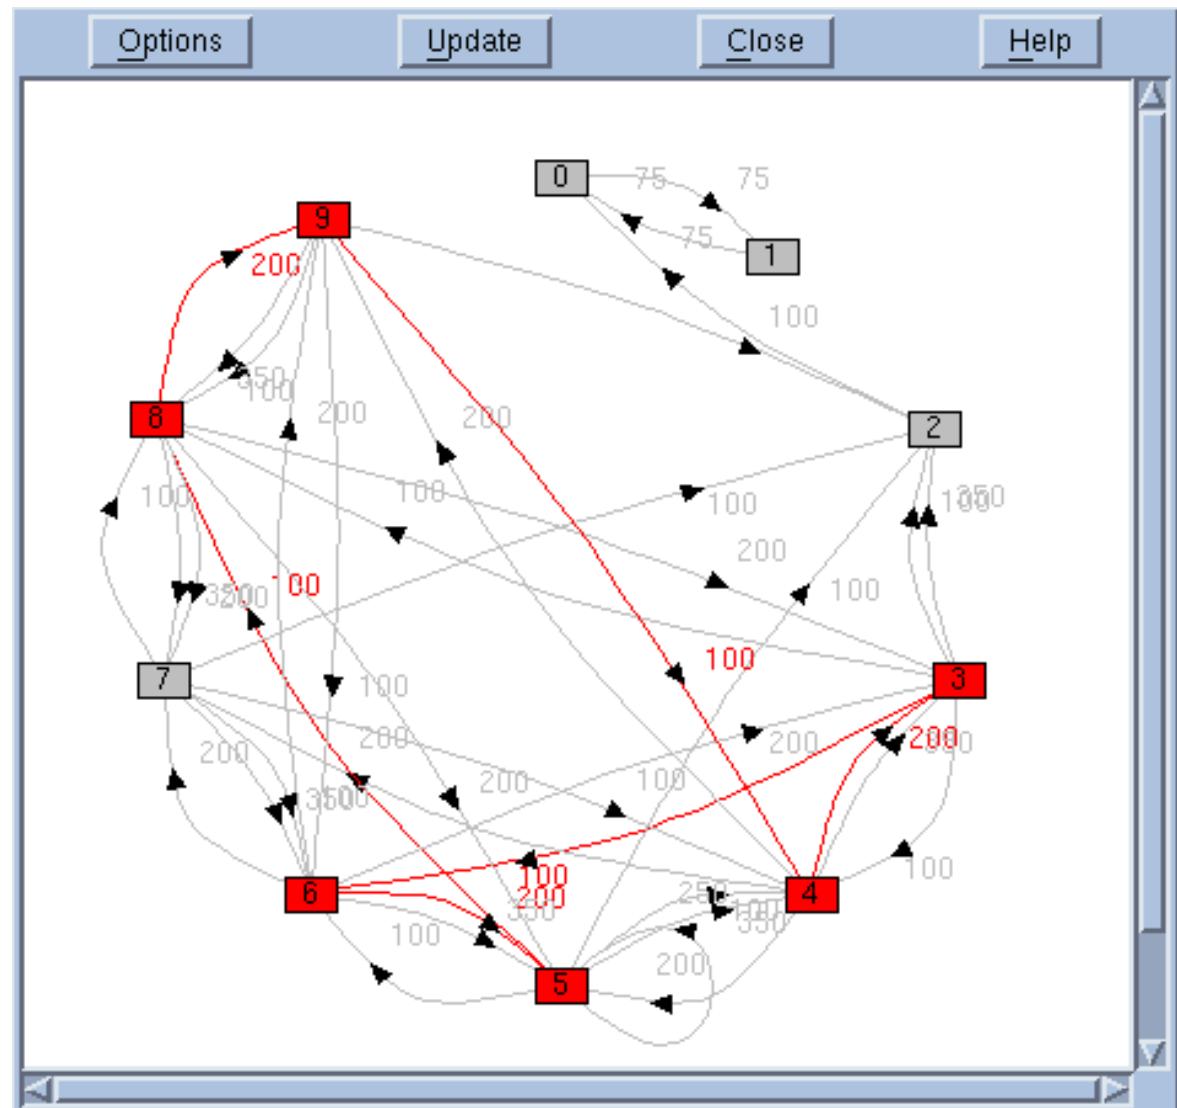
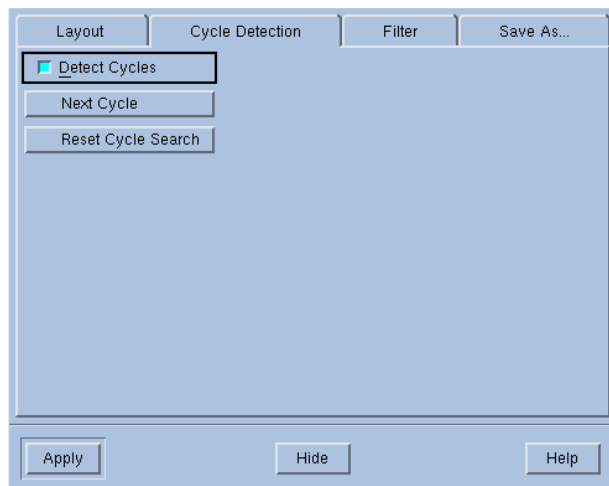
# Message Queue Graph

- Hangs & Deadlocks
- Pending Messages
  - Receives
  - Sends
  - Unexpected
- Inspect
  - Individual entries
- Patterns



# Message Queue Debugging

- **Filtering**
  - Tags
  - MPI Communicators
- **Cycle detection**
  - Find deadlocks



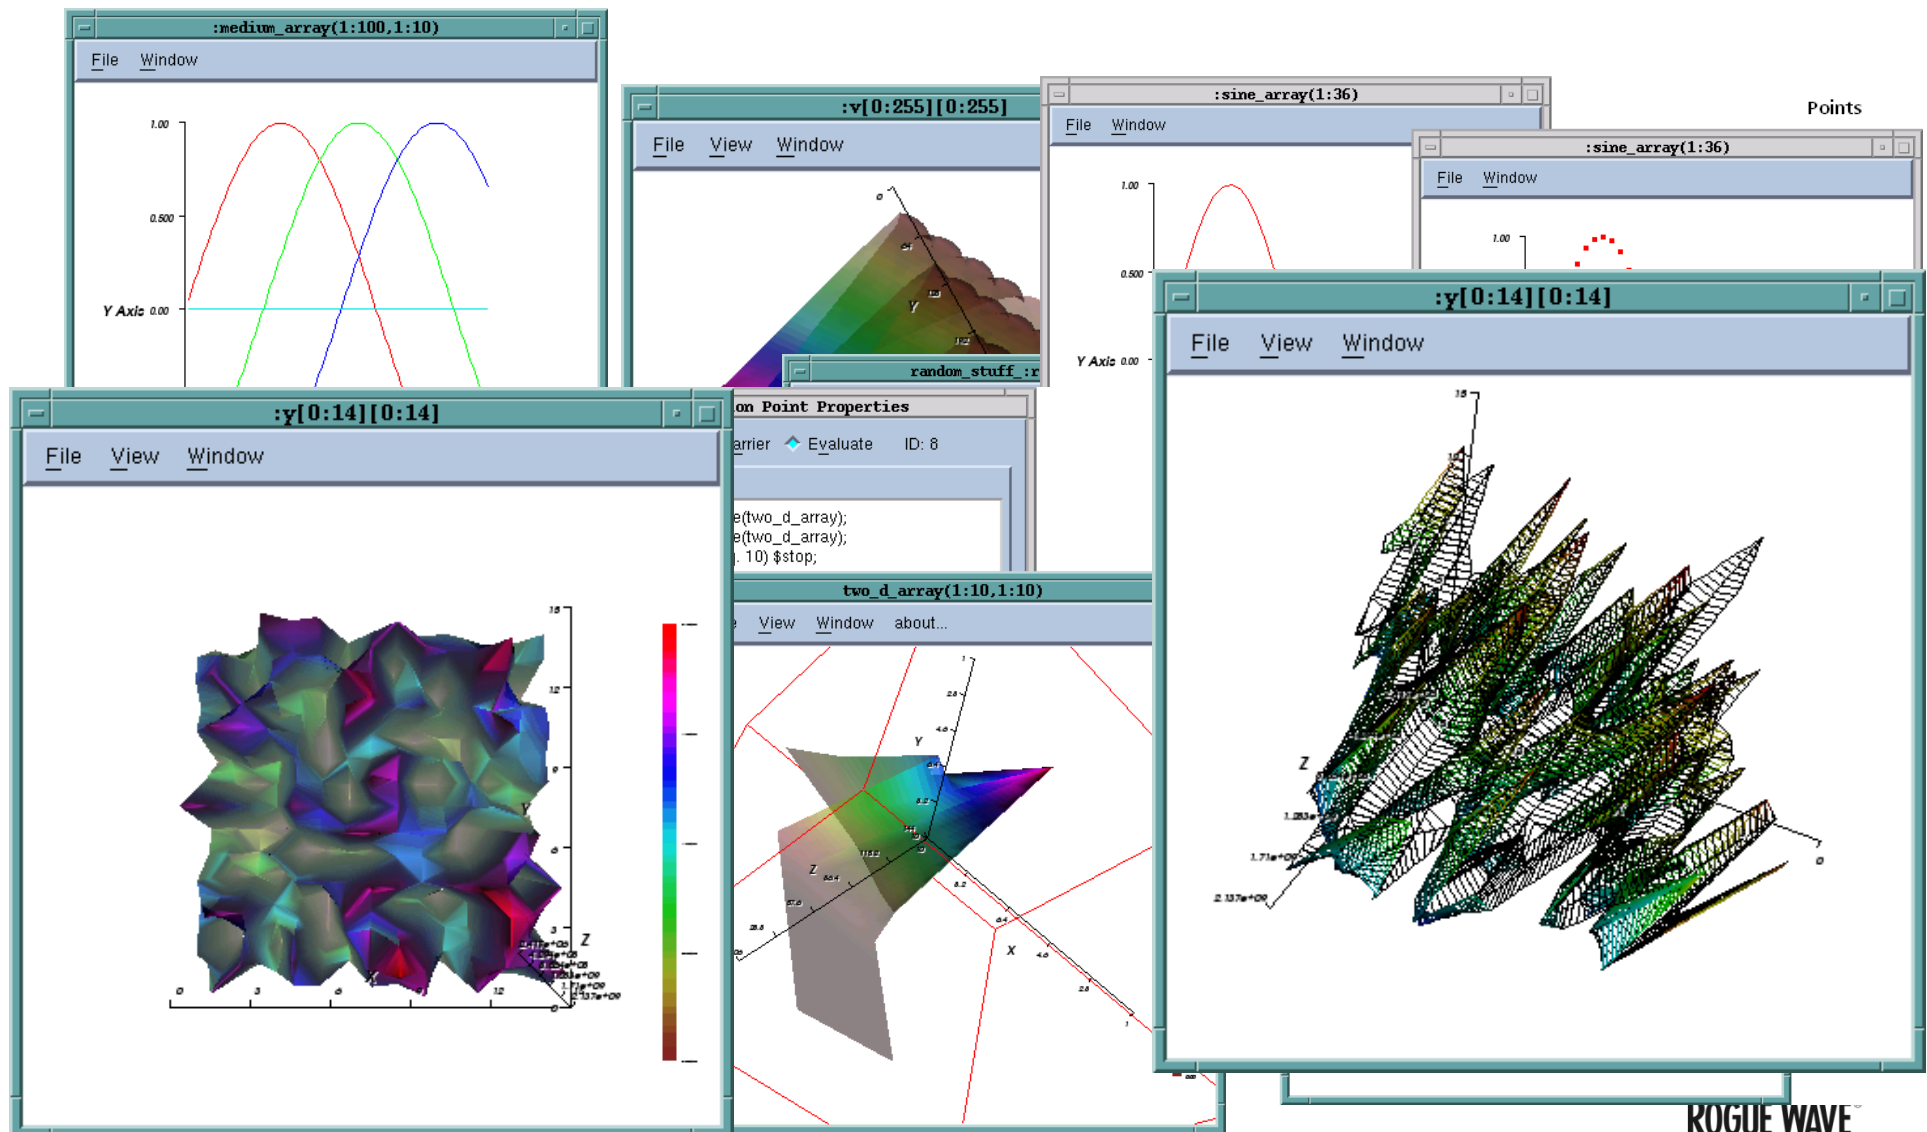


# Visualize

---

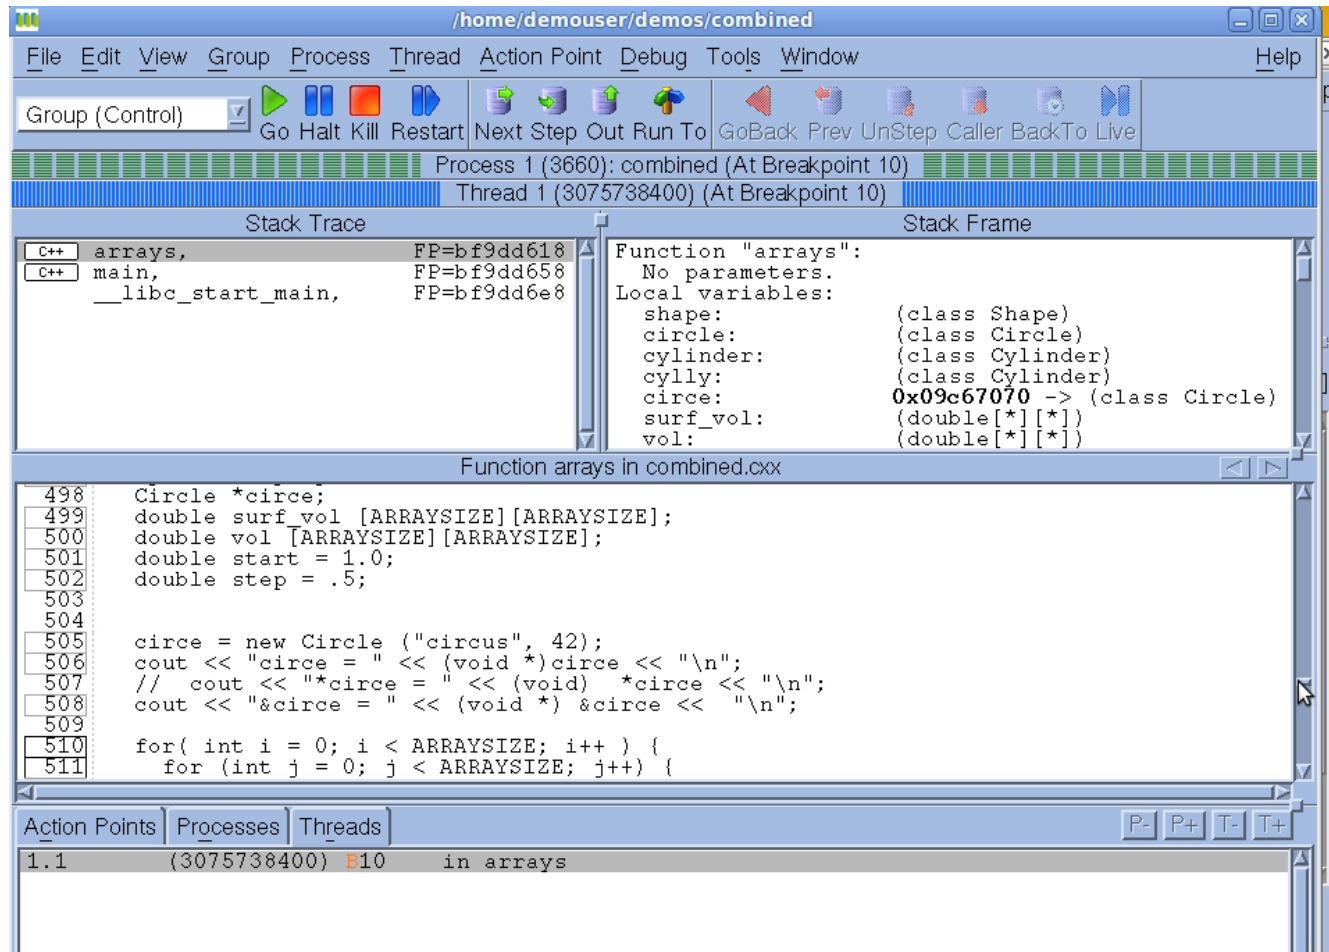


# Visualization



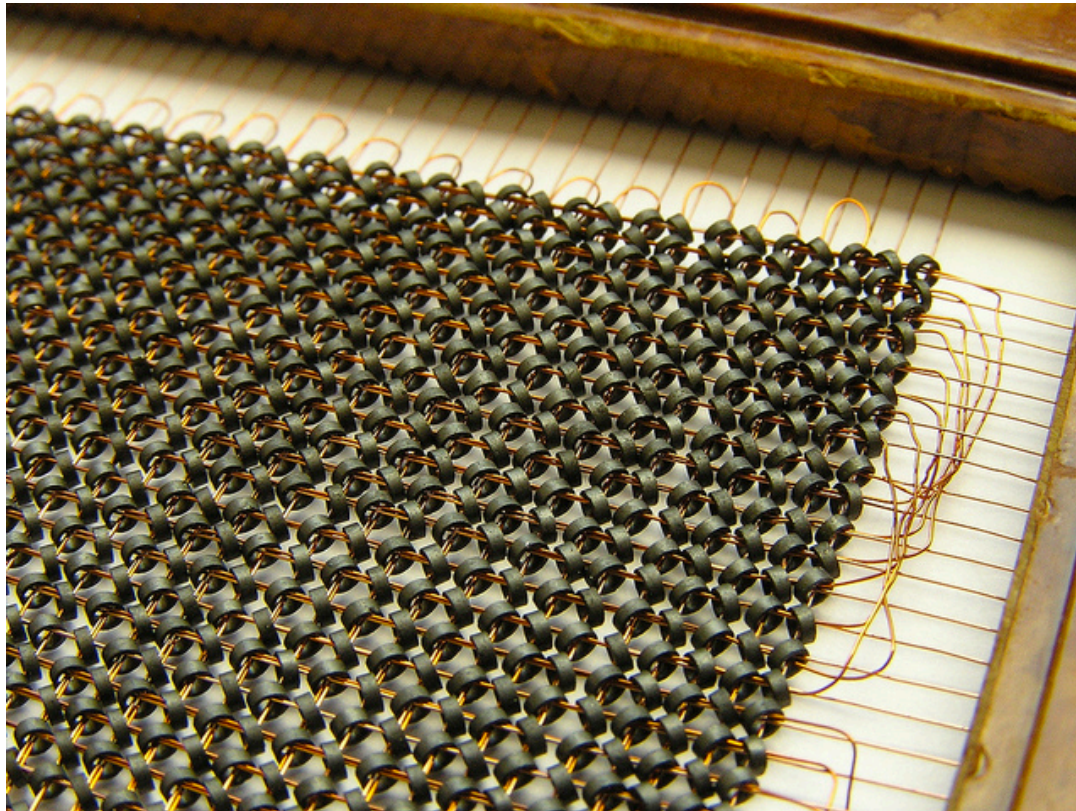
# Visualization

Get the big picture – Observe anomalies – Utilize Pattern recognition – Save time!



# ... And Don't Forget the Memory!

---



# MemoryScape

---

**Memory bugs often go undetected until the worst possible time**

- Symptoms often surface long after the actual damage is done
- Some only surface after hours or even days of operation
- In many cases, the programs affected are “innocent bystanders”

**MemoryScape: Fully Integrated in TotalView**

**No Source Code or Binary Instrumentation**

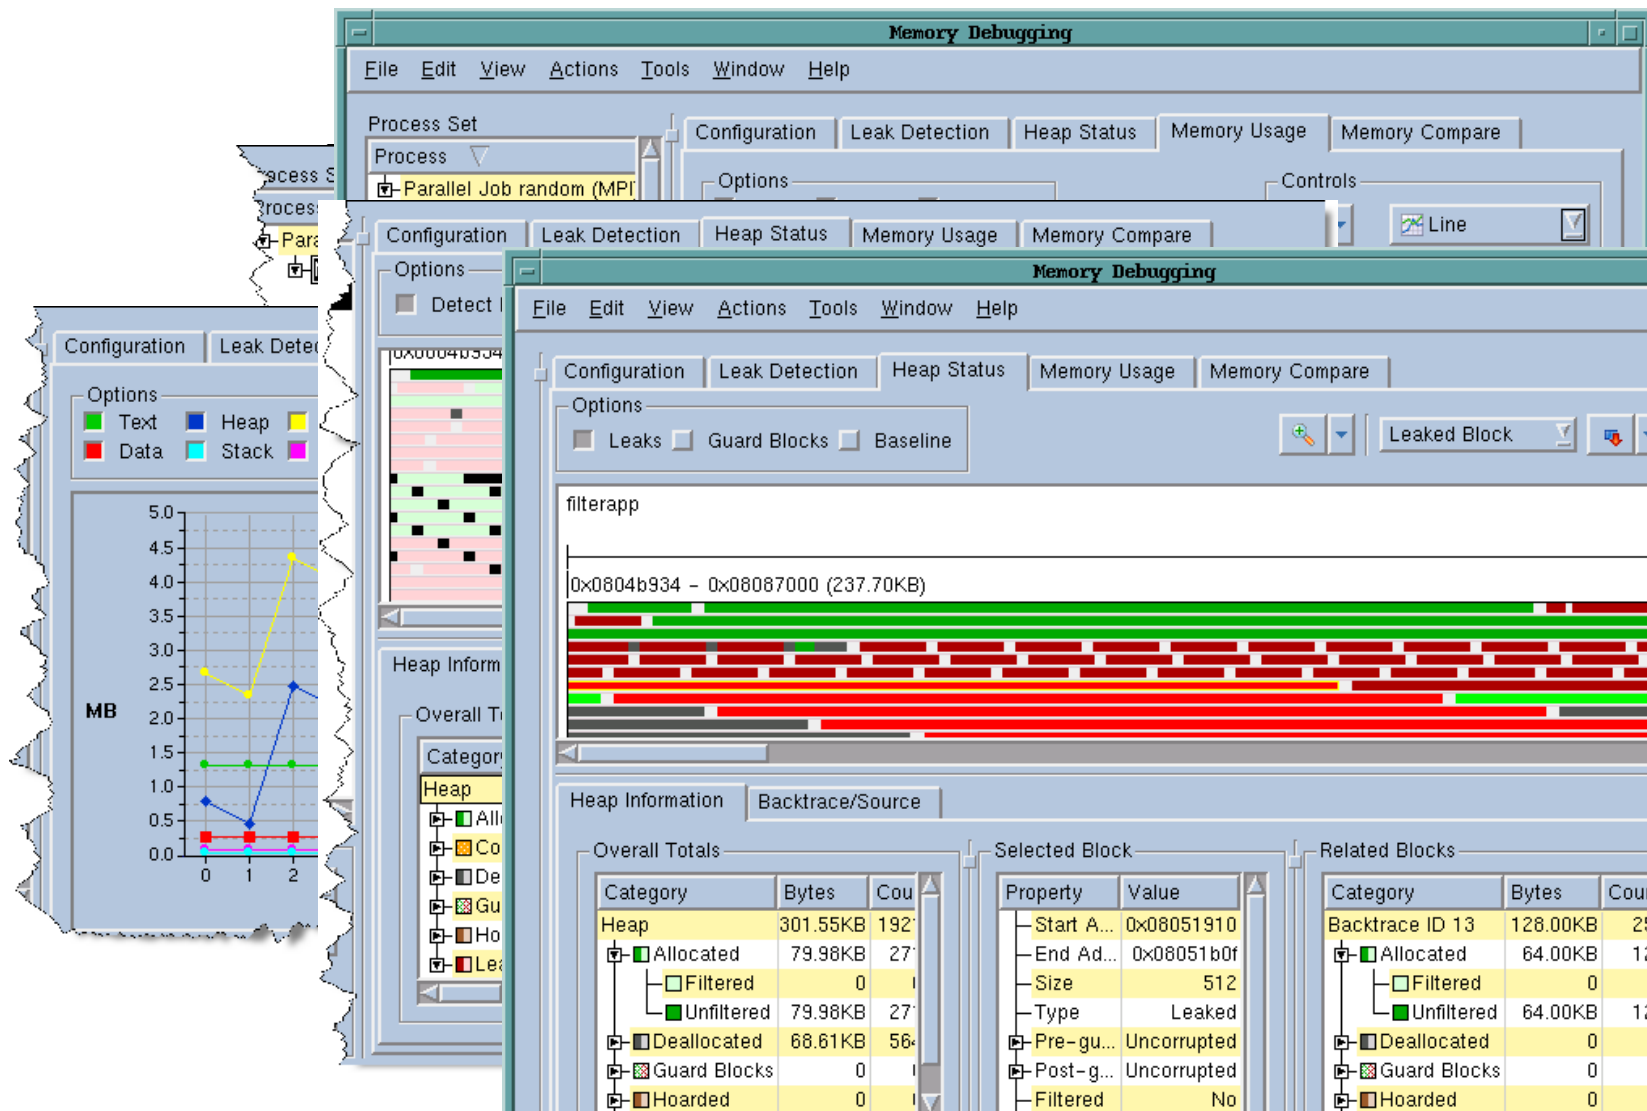
- Use it with your existing builds
- Programs run nearly full speed
- Low performance overhead
- Low memory overhead • Efficient memory usage

**MemoryScape Feature Highlights**

- Automatic allocation problem detection
- Heap Graphical View
- Leak detection
- Block painting
- Dangling pointer detection
- Deallocation/reallocation notification
- Memory Corruption Detection - Guard Blocks
- Memory Hoarding
- Memory Comparisons between processes
- Collaboration features



# MemoryScape



# What's Coming

---

- **Increased Scalability**
  - Leveraging TotalView's Architecture
  - Efficient Use of Cluster Resources
    - Extremely light weight debug agents; Minimal memory footprint
    - More space on the compute nodes for user application code
  - Tree-Based Overlay Network
    - Broadcast of Operations; Aggregation of Events and Data
- **Advanced User Interface**
  - New GUI Framework
  - Changes focused on extreme scale debugging
- **CUDA 4.1 now; 4.2 and 5.0 this year**
- **Replay Enhancements**
  - Record on Demand (in Beta)
  - Replay Debug from Core File
- **OpenACC Support**
- **Intel MIC Support**
  - Come see a demo at ISC '12



# Developing for Parallel Architectures



**TotalView®**

- **Code debugging**
  - **Highly scalable interactive GUI debugger**
    - Easy to use -- without sacrificing detail that users need to debug
    - Used from workstations to the largest supercomputers
  - **Powerful features for debugging multi-threaded, multi-process, and MPI parallel programs**
  - **Compatible with wide variety of compilers across several platforms and operating systems**
- **Memory Debugging**
  - **Parallel memory analysis and error detection**
    - Intuitive for both intensive and infrequent users
  - **Easily integrated into the validation process**
- **Reverse Debugging**
  - **Parallel record and deterministic replay within TotalView**
    - Users can run their program “backwards” to find bugs
  - **Allows straightforward resolution of otherwise stochastic bugs**
- **GPU CUDA Debugging**
  - **Full Hybrid Architecture Support**
  - **Asynchronous Warp Control**
  - **Multi-Device and MPI Support**

*Thanks!*



Developing parallel, data-intensive applications is hard.  
**We make it easier.**

[www.roguewave.com](http://www.roguewave.com)

